The USGS utilizes trade studies and reviews for our own business purposes but does not endorse any one vendor or product.

Architecture and Technology Program

February 2004

Offline Archive Media Trade Study

Offline Archive Media Trade Study

February 25, 2004

By SAIC

Prepared by:

Tom Bodoh Date
Principal Systems Engineer
SAIC

Reviewed by:

Tony Butzer Date
Architecture & Technology Task Lead
SAIC

Reviewed by:

Ken Gacke Date
Principal Systems Engineer
SAIC

Reviewed by:

Cheryl Greenhagen Date
Principal Systems Engineer
SAIC

Reviewed by:

Al Engelbrecht Date
Principal Electronic Engineer
SGT

Reviewed by:

Kenneth Boettcher Date
Archive Manager
SAIC

Preface

This document contains the Offline Archive Media Trade Study prepared by SAIC for the USGS. The Trade Study presents the background, technical assessment, test results, and recommendations.

This Trade Study was prepared by

SAIC (Science Applications International Corporation)
USGS/EROS Data Center
Sioux Falls, SD 57198

Team members:
Tom Bodoh
Ken Boettcher
Ken Gacke
Cheryl Greenhagen
Al Engelbrecht
Tony Butzer

Abstract

This document is a trade study comparing offline digital archive storage technologies. The document compares and assesses several technologies and recommends which could be deployed as the next generation standard for the USGS at the EROS Data Center (EDC). Archives must regularly evolve to the next generation of digital archive technology and the technology chosen must remain reliable until the next migration. Note that this study is a revisit of a study completed in FY01 (Fiscal Year 2001) and revised in FY03.

Revision History

2/25/04

- Changed to allow for consideration of helical scan as long as certain performance criteria are met
- Added LTO2 as a current archive technology
- Added SAIT-1 and SuperDLT 600 as considered drives
- Replaced IBM 3590 with IBM 3592
- Removed LTO1 and SDLT 320 from the study
- All drives in the study are considered
- Increased the minimum specs for capacity and transfer rate
- Reworked cost scenarios, and reduced to three
- Removed transfer time scenarios
- Removed maintenance from cost scenarios
- Removed criteria showing multi-vendor availability as an advantage

Contents

1.0 Introduction

1.1 Purpose and Scope

This document provides an assessment of the options for the next generation of offline digital archive storage technology to be used for the Digital Archives of the USGS. The selected technology must be capable of safely retaining data until space, cost, and performance considerations would drive migration.

It is envisioned that within two years most or all of the USGS archive holdings will reside on nearline storage and will be backed by an offline master copy and possibly an offsite copy. The nearline copy is referred to as the working copy. There continues to be a need for offline storage for infrequently used working copies, as well as master and offsite copies where the working copy is stored nearline.

Note that LTO2 is the current archive media of choice at the USGS EROS Data Center. There is no current desire by the USGS to change technologies at this time and given the advantages of inter-generation compatibility in an offline archive environment, there will be a continued interest in “staying the course” with LTO technology. For these reasons, the main purpose of the study is to stay current on offline storage technology so that if LTO begins to falter as a technology, the USGS will have the information necessary to make a change.

1.2 Background

The USGS, Earth Resources Observation Systems (EROS) Data Center, located in Sioux Falls, SD, currently archives offline datasets using several technologies. In 1992, the TMACS (TM/MSS Archive Conversion System) system was deployed to transcribe Landsat archives from HDT (High Density Tape) to DCT (Digital Cassette Tape). Both HDT and DCT utilize large, expensive analog instrumentation drives, which require frame synchronization, driving the cost of transcribing Landsat HDTs to DCTs to exceed $1,000,000 for each generation of media. Note that DCT and HDT are not purely analog. Although the crucial IRIG (InteRange Instrumentation Group) data is stored in analog format, the image data is stored in digital format. Though the conversion from HDT has been completed, transcription from DCT to digital media will soon begin.

Locating, rehabilitating, and integrating HDT and DCT drives has been costly in terms of labor, parts, and vendor service costs. The ongoing maintenance costs for the HDT and DCT drives are excessive compared to drives such as SDLT and LTO since there is little industry experience and only a single vendor to support each brand of drive. The HDT and DCT drives in existence today number in the dozens, with the count decreasing each year as other users transition to digital media. Three new transcription systems were implemented in the past five years, two of them transcribing HDT media to computer compatible DLT 7000. The latest transcription system will transcribe HDT or DCT data to disk files, which will be copied to nearline and offline media.

The DLT 7000 drive was retired by Quantum several years ago, replaced by the DLT 8000. The DLT 8000 has not been widely accepted since the SuperDLT drives had already been pre-announced when the DLT 8000 was released. A USGS study of DLT 7000 errors revealed that they exhibit a greater percentage of data loss as compared to 3480, 3490 and 9840.

Table 1-1 summarizes the offline archive tape technologies currently in use:

Tape Drive Technology	Capacity	Transfer rate	Type
HDT	3.4 GB	10.6 MB/sec	Analog
3480	200 MB	2 MB/sec	Digital
3490	900 MB	2.7 MB/sec	Digital
DLT 7000	32 GB	4.7 MB/sec	Digital
DCT (Ampex DCRsI)	45 GB	12 MB/sec	Analog
SuperDLT 220	98.8 GB	8.1 MB/sec	Digital
HP LTO Ultrium 2	197 GB	23.8 MB/sec	Digital

Table 1-1 Past and Current Archive Technologies Used

In 2003, the USGS migrated more than 50,000 3480 and 3490 tapes to nearline storage, and 110 LTO2 tapes. This migration was performed over a period of 5.5 months, slowed by the handling of the large number of 3480/3490 tapes. This migration freed up enough library shelving to ensure that the library should never need to be expanded, and may in fact be reduced in size. Several other smaller migrations are continuing, with the intent to reduce the offline archive to as few tapes as possible.

The USGS has utilized SuperDLT 220 extensively for onsite and offsite backups. HDT, 3480/3490, and DCT have proven to be robust and high-performance for their time. As technology advances, as datasets grow, as media ages, and as USGS Digital Library space fills, the USGS must migrate data to newer, more physically compact, and higher performing storage technologies.

1.3 Data reliability

Since the foremost goal of an archive is data preservation, the primary criteria for the selection of the drive technology must be reliability. Several elements contribute to data reliability:

- The number of archival copies: The dependence on the master copy, and the level of risk rise when a working copy is not robust. Note that the master and working copies need not be on similar media, though generation and recovery of a working copy is simplified if the storage capacities are similar. All USGS archives must have both working and master copies, and an offsite copy is desirable. Note that a slightly less reliable drive technology can be used if there are a sufficient number of copies of the archive.

- The storage location and environment: This is a constant for all of the technologies assessed since any would be stored in a secure and climate-controlled environment.

- The composition of the media: Some media compositions last much longer than others, though all of the technologies in this study use similar long-lasting media compositions.

- Tape handling within the drive: This characteristic defines how a tape is handled by the drive, whether contact is made with the recording surface, how many passes are required to read or write an entire tape, and the complexity of the tape path.

- Error handling: The ideal drive minimizes data loss through CRC (Cyclic Redundancy Check) or other data recovery methods, and allows data to be read after skipping over an error. Though error detection upon write is required, additional attention to data recovery upon read is a higher priority since media degradation will lead to eventual read errors.

- Primary Market: This criterion describes the target market of a drive, and the characteristics of drives within that market. A drive targeted to the backup market would be designed for write many/read rarely, therefore more emphasis is placed on detecting errors upon write. A drive targeted to the archival market would be designed for write once/read few and more emphasis is placed on detecting and correcting errors upon read. A drive targeted to the Enterprise market would be designed for write many/read many and equal emphasis is placed on detecting errors upon read and write. Ideally, all archives would be written to a drive designed for the archive market, but none are currently available. Most vendors would argue that their products are archive devices, but if forced to choose their primary market no vendor would choose the very limited archive market over the lucrative backup market. With proper handling and multiple copies, any of the technologies evaluated in this report could be deployed for archive use.

Primary Market	Reliability	Usage	Driving Design Factors
Backup	Moderate	Write many, read rarely	Cost, capacity, speed
Enterprise	High	Write many, read many	Designed for continual use, often with robotics
Archive	High	Write once, read few	Long term reliability

Table 1-2 Tape Drive Markets and Characteristics

The reliability of a long-term archive technology relates primarily to the long-term viability of the recorded media. Since it is wise to implement a technology early enough in its life cycle that drives can be kept viable through the lifetime of a given media (or replaced with newer backward-compatible models), a definitive leader in reliability is difficult to determine except in retrospect. This study bases the reliability assessment on past experience with the vendor and their products, on specifications, on the experiences of others, or experience gained from benchmarking.

StorageTek 9940 uses serpentine recording but uses many fewer passes than either LTO or SDLT. In addition, 9940 drives do not touch the recording surface, and redundant servo tracks are provided. Experience with 3480, 3490, 9840, 9940A and 9940B has shown StorageTek products to be very reliable. The StorageTek D3 helical scan drive was problematic and was discontinued quickly. On two occasions, 9840 tapes that encountered unrecoverable errors were sent to StorageTek for recovery. One tape was recovered, but the other was unrecoverable due to cartridge contamination.

1.4 Technologies selected for consideration

The criteria used in determining which technologies should be considered were:

1. The technology must be currently available and shipping in order to be considered in the final analysis. It also must be the latest drive in the line.

2. The technology must hold at least 190 GB of data.

3. The technology must have a write transfer rate of at least 20 MB/sec.

4. The technology must use a media that can remain readable for at least 10 years in a controlled environment. The lifetime of 10 years was selected since it is the longest that a media technology would conceivably be used before space and transfer rate concerns would dictate a move to a new technology.

5. The technology must not be hampered by a poor reliability history. Though helical scan technologies such as 8mm, 4mm, DAT and D3 have proven unreliable in the past for archive purposes, newer helical scan technologies may have addressed earlier issues.

The currently available drive technologies selected for consideration are:

1. StorageTek 9940B

2. HP LTO2 (Linear Tape Open)

3. IBM LTO2 (Linear Tape Open)

4. Quantum SuperDLT 600

5. Sony SAIT-1

6. IBM (International Business Machines) 3592

1.5 Dismissed technologies

The following technologies were dismissed from further analysis or consideration for the reasons listed.

1.5.1 CD-ROM, DLT 8000, QIC, and Erasable Optical (EO)

This category includes technologies that are low capacity, low performance, or aged. All of these products have been available for some time, but can immediately be dismissed based on obvious limitations in performance, capacity, or reliability. These products are clearly not a good fit for large digital archives.

1.5.2 Exabyte VXA2 and Mammoth 2

Exabyte has evolved its early helical scan technology into two product lines: VXA2 with a native capacity of 80 GB a native transfer rate of 6 MB/sec and the Mammoth 2 with a native capacity of 60 GB and a native transfer rate of 12 MB/sec. These technologies are based on consumer-grade cartridge and drive technologies. While media costs are low, transfer rates are low and the USGS experience with consumer-grade storage technology has shown that this technology cannot withstand the rigors of long-term archive.

1.5.3 DVD

DVD (Digital Video Disc) seems promising from the standpoint of longevity of the media. However, low capacity per media, low transfer rates, lack of media protection, no single standard, and high media costs add up to a product that simply would not work for high volume archival use.

1.5.4 Other technologies

Several high capacity optical disk technologies have been in the development phase for the past few years. Of the 100+ GB technology proposals that have appeared in trade journals and at conferences, to date none are shipping products.

Other high-tech examples of future technologies such as holographic storage or bio-storage will not mature for several years.

2.0 Technical Assessment

2.1 Analysis

StorageTek 9940B:

Advantages:

- USGS experience with 9940 drives at EDC has shown them to be extremely reliable.
- The 9940B has ‘wider’ tracks (16 tracks per pass instead of 8) which reduces serpentine passes.
- The 9940B has air bearings that allow the tape to float past the head without contact.
- 9940 is targeted to the Enterprise Storage market where data viability, speed, and capacity are more important than cost.
- 9940 was designed as a robust storage media, with the tape cartridge and drive built to withstand constant and/or frequent use in a robotic environment. The 9940 drives are compatible with the USGS StorageTek silos and excel in a robotic environment due to their durability.
- The USGS offline Digital Library shelving and tape carriers used for 3480/3490 accommodate 9940 cartridges.
- 9940 drives provide drive statistics such as servo errors, bytes read/written, I/O retries, and permanent errors.

Disadvantages:

- StorageTek has indicated that the 9940B is the last of the 9940 series, as they have reached the limits of metal particle technology in the 9940 design. They are working on a new product (9950?) with the first new drive at 500 GB native capacity shipping in 2004 or 2005. The drive will use new media and it is unclear whether this drive would be backward compatible.
- The drives are relatively expensive, as compared to SDLT, SAIT or LTO.

Notes:

- The usable capacity may vary between cartridges. The USGS attained a capacity of 193.03 GB per tape. CERN (Conseil European pour la Recherché Nucleaire) was able to write 208 GB per cartridge on all 10 tapes they tested (http://cscct.home.cern.ch/cscct/T9940B.ppt).
- StorageTek is the sole manufacturer of 9940B. This may be seen as an advantage given the incompatibilities experienced between LTO drives manufactured by HP and IBM
- While the projected follow-on 9950 will take different media, it is anticipated that StorageTek will continue the tradition of using the same physical cartridge dimension so that existing robotic libraries can accept the new media without modification. This should also ensure compatibility with offline shelving.

LTO2:

Advantages:

- LTO has enjoyed phenomenal growth from the day of release, continuing through the recent slowdown in the IT sector despite contraction of the tape industry as a whole.

- LTO has a 67% market share, with 250,000 drives installed worldwide, compared to 150,000 SDLT drives.

- Backward read/write compatible with LTO1. This means that the LTO2 drive can read and write LTO1 cartridges in the LTO1 density. All future drives are slated to be able to read any previous generation of tape.

Disadvantages:

- LTO is targeted to the backup market where speed, capacity, and cost are more important than long-term viability of the data. Since backup tapes are write-many/read-rarely, errors would likely show up in a write pass where they can be worked around (rewrites) or the media discarded.

- Repeated end-to-end use of a tape would be a concern since one end-to-end read/write would incur 64 passes.

- LTO was co-developed by Seagate, IBM, and HP (Hewlett Packard). This type of deployment makes it possible for each vendor to interpret the specifications differently, and to design drives which may have incompatibilities. EDC has observed two incompatibility problems between HP and IBM: tapes written to EOT on the IBM cannot be read on the HP, and tapes written on the HP reads at less than half speed on the IBM.

- LTO was designed as a moderate usage storage media, with the tape cartridge and drive not built to withstand constant or frequent use. Although STK added the capability for their large silo to handle both LTO and 9940/9840, the robotic arm had to be slowed down since the thin shell of the LTO could not take the grip pressure necessary to keep cartridges from flying out of the gripper when the arm is at full speed.

Notes:

- IBM and HP are shipping LTO2 drives. EDC has three HP LTO2 drives, and one IBM LTO2 drive. The IBM far outperforms the HP, but incompatibility issues have been encountered.

- The USGS AVHRR archives were recently copied from 50,000+ 3480 and 3490 tapes to just 110 LTO2 tapes, utilizing HP drives. No errors were encountered.

- The third and fourth generations of LTO Ultrium have been projected but not scheduled. LTO3 and LTO4 will have native capacities of 400 and 800 GB and native transfer rates of 80 and 160 MB/sec. The LTO consortium does not estimate dates for future products, but industry estimates are for LTO3 to be released in 2004 or 2005.

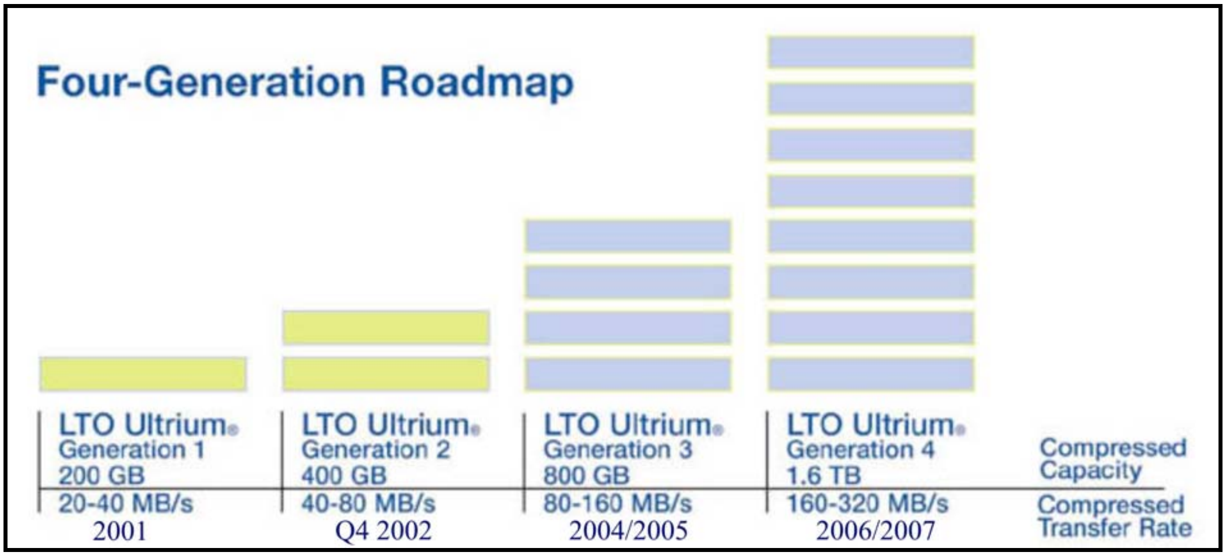

Figure 2-1 LTO Roadmap (with 2:1 compression)

SDLT 600:

Advantages:

- SDLT has enjoyed very wide market saturation due to its history in the backup market, although LTO has overtaken SDLT.
- Capacity exceeds 9940B and LTO2 at 300GB advertised native capacity.

Disadvantages:

- SDLT 600 uses different media than SDLT 220 and 320, ensuring that media costs will be significantly higher until market saturation drives the price down.
- SDLT is targeted to the backup market where speed, capacity, and cost are more important than long-term viability of the data. Since backup tapes are write-many/read-rarely, errors would likely show up in a write pass where they can be worked around (rewrites) or the media discarded.
- Repeated end-to-end use of a tape would be a concern since one end-to-end read/write would incur 80 passes.
- SDLT was designed as a moderate usage storage media, with the tape cartridge and drive not built to withstand constant or frequent use.

Notes:

- The product roadmap below calls for a drive capable of 600 GB at 64 MB/sec by the second half of 2004, and 1.2 TB (Terabytes) at 100 MB/sec in 2006.

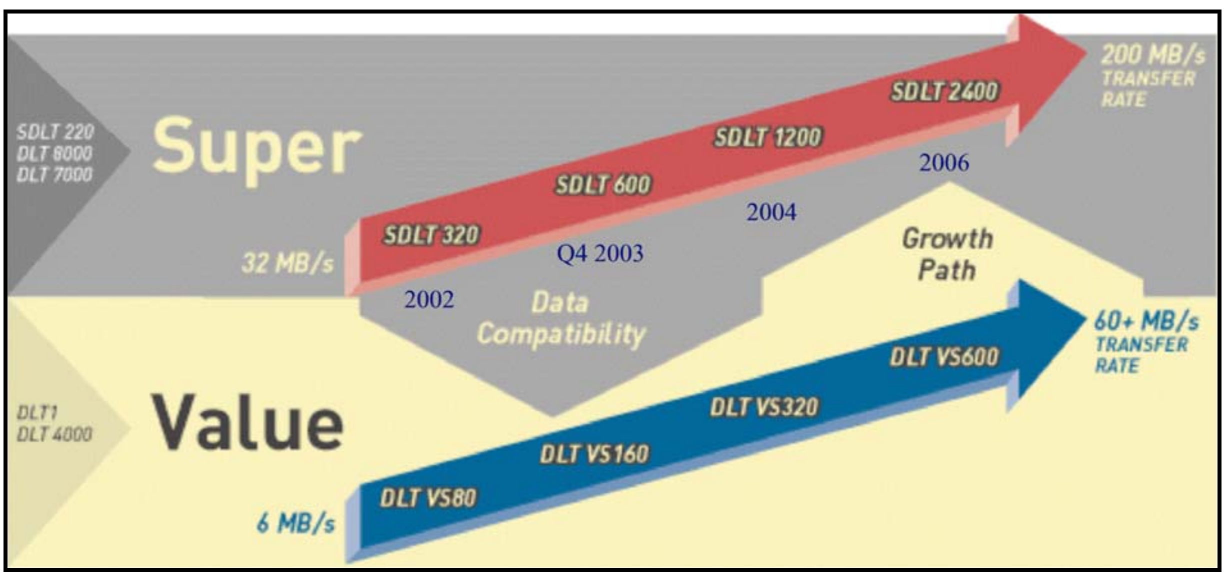

IBM 3592:

Advantages:

- Based on the very reliable 3480, 3490 and 3590.
- Excellent transfer rate and capacity.
- The 3592 line was designed as a robust storage media, with the tape cartridge and drive built to withstand constant and/or frequent use.

Disadvantages:

- The 3592 takes different media than the 3590 line.

Notes:

- The exact tape drive model number is 3592-J1A.
- IBM has a product roadmap showing future 3590 drives having capacities up to 1000 GB native, and transfer rates reaching 160 MB/sec. No timeline is provided but a prototype of the 1000 GB was running as of May 2002. The 3590 series seems to have evolved into the 3592.
- Although the 3592 is a follow-on to the 3590, the cartridges are not compatible.

Sony SAIT:

Advantages:

- Highest native capacity.
- Excellent transfer rate.
- Single pass (vs. serpentine) reduces tape movement, though head contact may actually be higher due to helical scan technology.

Disadvantages:

- Helical scan technologies have proven unreliable in the past due to complex drive path, high and constant head contact, tendency to destroy tapes, poor transfer rates, and extremely poor start/stop/repositioning times. Time will tell whether this heritage has been overcome.

Notes:

- It is unclear whether major robotic manufacturers can easily adapt current robots to accept SAIT in addition to SDLT, LTO, and 9940. Their ability to adapt depends on cartridge dimensions, cartridge rigidity, and barcode label size/location. Additionally, robotic vendors such as STK may not be willing to support this drive for competitive reasons, though they do support and sell SDLT and LTO technology in their current line of robots.
- Sony is projecting three more generations of SAIT, with SAIT-4 having a native capacity of 4 TB. No timelines or transfer rate projections are given.
- LP-DAAC staff experienced serious problems with other types of Sony tape drives, but Sony has been very responsive in resolving the issues. This may not be indicative of typical support, as the DAAC is a very large customer.

3.0 Tables

3.1 Design criteria

The design criteria and target market of a drive are interrelated. Drive technologies such as LTO, SAIT and SDLT are targeted to the backup market as demonstrated by their marketing. The 9940B and 3592 are targeted to the Enterprise (data center) market.

A drive targeted to the backup market is designed for write many/read rarely and more emphasis is placed on detecting errors upon write. Backup drives are typically built for speed, capacity, and low cost. A drive targeted to the Enterprise market is designed for write many/read many use in a robotic library or auto-stacker and equal emphasis is placed on detecting errors upon read and write. A drive targeted to the archival market is designed for write once/read few and more emphasis is placed on detecting and correcting errors upon read.

Enterprise and Archival drives are typically built for reliability, with speed and cost a secondary factor. All drives attempt error detection and recovery upon both read and write, but an archival drive design typically places more importance on read data recovery since data may no longer be available. Conversely, a backup drive places more importance on write error detection since the data is still available and can be easily rewritten.

The formula used to rank Design Criteria was:

((100-passes)/10)+
(error factor/2)+
(construction 3=moderate, 5=high usage)+
(head contact 3=yes, 5=no)
/ 2.54 (to adjust the highest rank to 10)

Technology	Serpentine tracks/ Passes	Target Market	Tape Composition	Uncorrected Error Rate	Cartridge Construction Rating	Head Contact	Ranking
STK 9940B	576/36	Enterprise	Advanced MP	$1x10^{18}$	High usage	No contact	10.0
IBM 3592	512/64	Enterprise	Advanced MP	$1x10^{17}$	High usage	Contact	7.9
HP/IBM LTO2	512/64	Backup	Metal Particle	$1x10^{17}$	Moderate usage	Contact	7.1
Sony SAIT-1	80*	Backup	Advanced ME	$1x10^{17}$	Moderate usage	Contact	6.5
SDLT 600	640/80	Backup	Advanced MP	$1x10^{17}$	Moderate usage	Contact	6.5

Table 3-1 Design Criteria and Target Market

*Prior experience with helical scan dictates caution. Helical scan cannot be rated for serpentine passes, but the equivalent is the helical passes per inch which are believed to be high (but unknown). For this reason SAIT was assigned the same number of passes as the SDLT 600.

3.2 Transfer Rate

Transfer rate is important since it dictates how quickly the migration of an archive data set may be completed, and how fast a production system may generate products from the archive media. Although 20 MB/sec is the minimum, it is desired to attain 30 MB/sec. Since much of the data archived at the USGS is not compressible, all transfer rates are native (uncompressed).

Note that tests performed by EDC did not utilize imbedded tape marks, improving performance but complicating the retrieval of individual files. This is acceptable for an archive tape that is not used as a working copy. Note that the source of the test results also applies to capacities in table 3-3.

The ranking was determined by adding the actual/estimated read and write rates for each drive, setting the ranking for the fastest drive to 10, then ranking the others against the leader. As an example, a drive having half of the total read/write transfer rate of the leader would be ranked 5.

Tape Drive Technology	Advertised Native Rate	Source Of Test Results	Actual/estimated Native Write Transfer Rate	% Of Adv.	Actual/estimated Native Read Transfer Rate	% Of Adv.	Ranking
SDLT 600	36 MB/sec	Estimate	33.26 MB/sec est.	92.4%	36.00 MB/sec est.	100%	10.0
Sony SAIT-1	30 MB/sec	Vendor	31.03 MB/sec	100+%	31.00 MB/sec	100+%	9.0
IBM LTO 2	35 MB/sec	EDC	32.70 MB/sec	93.4%	29.08 MB/sec	83.1%	8.9
HP LTO 2	30 MB/sec	EDC	28.90 MB/sec	96.3%	29.12 MB/sec	97.0%	8.4
STK 9940B	30 MB/sec	EDC	27.65 MB/sec	92.2%	29.28 MB/sec	97.6%	8.2
IBM 3592	40 MB/sec	Estimate	27.64 MB/sec est.	69%	20.00 MB/sec est.	50%	6.9

Table 3-2 Transfer Rates

3.3 Capacity

A secondary requirement is to conserve archive rack space and reduce tape handling by increasing per media capacity. The current archive media of choice at the USGS is LTO 2 at 197 GB per tape. The new minimum capacity requirement is 190 GB, with 250 GB or more desired. All of the reviewed technologies meet the 190 GB requirement. Since much of the data archived is not compressible, all capacities are native (uncompressed).

The ratings were determined by calculating each as the percentage of the highest capacity drive, on a scale of 1 to 10, with the highest capacity as a 10. The source of the capacity ratings are as noted in table 3-2 above.

Tape Drive Technology	Advertised Native Capacity	Measured/Estimated Native Capacity	% Of Advertised Capacity	Ranking
Sony SAIT-1	500 GB	475.00 GB estimated	95.0% estimated	10.0
SDLT 600	300 GB	286.80 GB estimated	95.6% estimated	6.8
IBM 3592	300 GB	285.00 GB estimated	95.0% estimated	6.0
HP/IBM LTO 2	200 GB	197.00 GB	98.5%	4.1
STK 9940B	200 GB	193.03 GB	96.5%	4.1

Table 3-3 Storage Capacities

3.4 Cost Analysis

Table 3-4 shows the relative drive and media costs, maintenance costs, and the cost per Terabyte for media. Rankings were established by setting the cheapest (drive, maintenance, media) to 10 then rating each of the others against the lowest cost. Maintenance is estimated at 15% of purchase price, which is an industry trend. STK maintenance is actual, based on the most recent support contract. Maintenance on the cheaper drives would entail utilizing a spare along with depot repair.

Technology	Drive $/each	Annual Maintenance	Media $/each	Media $/TB	Ranking Drive Cost	Ranking Maint Cost	Ranking Media Cost
STK 9940B	$28,000	$1,705	$72	$373	1.4	3.4	10.0
HP LTO 2	$3,836	$575	$75	$380	10.0	10.0	9.8
IBM LTO 2	$5,798	$869	$75	$380	6.6	6.6	9.8
SDLT 600	$4,995	$749	$110	$384	7.7	7.7	9.7
SONY SAIT-1	$9,979	$1,496	$190	$400	3.8	3.8	9.3
IBM 3592	$32,000	$2,424	$135	$473	1.2	1.2	7.9

Table 3-4 Drive, Maintenance and Media Costs

3.5 Scenarios

Table 3-5 shows the total drive and media cost for three scenarios. These scenarios presume that each dataset or project stands on their own, but pooling resources for multiple datasets can mitigate cost. Note that prices are expected to drop considerably within six months after product introduction. Rankings were established by setting the cheapest to 10 then rating each of the others against the lowest cost.

Technology	100 TB 2 drives	200 TB 4 drives	400 TB 6 drives	Ranking
HP LTO 2	$45,672	$91,344	$175,016	10.0
SDLT 600	$48,390	$96,780	$183,570	9.4
IBM LTO 2	$49,596	$99,192	$186,788	9.2
Sony SAIT-1	$59,958	$119,916	$219,874	7.6
STK 9940B	$97,300	$194,600	$329,200	4.7
IBM 3592	$111,300	$222,600	$381,200	4.1

Table 3-5 Scenario Costs (drives, media)
(Yellow indicates scenario ranked)

3.6 Vendor analyses

Table 3-6 is intended to provide an analysis of each company and the stability of each technology. All seem to be established and stable companies, and this rating should in no way be viewed as a market analysis. When selecting an archive technology, it only makes sense to look at company and product history. The longevity rankings were determined by the following formula:

(company age/10) +
(technology age))
* .55 (to adjust the highest rank to 10)

Company	Technology	Years in business	Technology age in years	Longevity Ranking
IBM	3592 (3590)	93 (1911)	9 (1995)	10.0
Quantum	SDLT (DLT)	24 (1980)	15 (1989)	9.6
Sony	SAIT (AIT)	58 (1946)	8 (1996)	7.6
IBM	LTO	93 (1911)	4 (2000)	7.3
HP	LTO	65 (1939)	4 (2000)	5.8
StorageTek	9940	35 (1969)	4 (2000)	4.1

Table 3-6 Vendor Analyses

3.7 Drive compatibility

Table 3-7 shows the level of inter-generation drive compatibility as well as the future drives planned. The column "Previous Generations Read" indicates how many previous generations are read by the generation indicated. The column "Future Generations Planned" indicates the number of generations planned in the current drive family, following the drive indicated. The column "Future Generations Compatible" indicates how many future generations are announced to be compatible with the drive indicated. The ranking was determined by the following formula:

(Previous Generations Read + Future Generations Planned + Future Generations Compatible) * 1.67 (to adjust the highest rank to 10)

Technology	Previous Generations Read	Future Generations Planned	Future Generations Compatible	Ranking
SDLT 600	2	2	2	10.0
Sony SAIT-1	0	3	3	10.0
IBM 3592	0	2	2	6.7
HP/IBM LTO2	1	2	2	8.3
STK 9940B	1	0	0	1.7

Table 3-7 Drive Compatibility

3.8 Ranking summary

The ranking summary provides a quick reference to the rankings. Drive and maintenance cost have been combined into one column since the scores are the same. For STK the drive and maintenance cost ratings differed, so the two rankings were averaged.

Drive	Design Criteria	Transfer Rate	Capacity	Drive/ Maint Cost	Media Cost	Scenario Cost	Vendor And drive Longevity	Drive Compat.
STK 9940B	10.0	8.2	4.1	2.4	10.0	4.7	4.1	1.7
Sony SAIT-1	6.5	9.0	10.0	3.8	9.3	7.6	7.6	10.0
HP LTO2	7.1	8.4	4.1	10.0	9.8	10.0	5.8	8.3
IBM LTO2	7.1	8.9	4.1	6.6	9.8	9.2	7.3	8.3
SDLT 600	6.5	10.0	6.8	7.7	9.7	9.4	9.6	10.0
IBM 3592	7.9	6.9	6.0	1.2	7.9	4.1	10.0	6.7

Table 3-8 Ranking Summaries
(Blue indicates the highest ranking in category)

4.0 Conclusions and Recommendations for USGS Offline Archiving Requirements

4.1 Weighted Decision Matrix

The following table provides a weighted analysis of the drives considered. The criteria emphasize the importance of traits contributing to data preservation. The USGS made the final decision regarding which criteria to use and the relative weighting of the criteria. The columns in green are relative ratings for each technology. The columns in yellow are calculated by multiplying the relative weight by the relative rating. The following describe each criterion:

- Design Criteria (Reliability of media): This criterion describes the ability of the media to remain readable over time. Included in this criterion is the number of passes per full-tape read or write, cartridge construction, uncorrected bit error rate (BER) and amount of head contact. (See table 3-1)
- Capacity: This criterion describes the measured or estimated capacity per cartridge, which is typically less than the advertised capacity. (See table 3-3)
- Media cost/TB: This criterion is a rating of the relative cost per Terabyte for media using the measured or estimated capacity rather than advertised capacity. (See table 3-4)
- Compatibility: This criterion describes the likelihood that the drive technology will continue to evolve and the extent to which future drives will have backward read capability. This will give an indication of the ability to maintain drives that can read an aging archive. (See table 3-7)
- Transfer rate: This criterion describes the aggregate read and write transfer rate, which is typically less than the advertised transfer rate. (See table 3-2)
- Drive cost: This criterion is the rating of relative cost of each drive at the lowest currently available price. (See table 3-4)
- Vendor analyses: This criterion is the rating of the viability of the vendor and technology. (See table 3-6)
- Scenario cost: This criterion is the rating of the cost of scenario #1. This includes media cost and drive cost. The measured or estimated capacity is used rather than advertised capacity. (See table 3-5)

Note that in the decision matrix spreadsheet below, not all criteria have been selected for the final analysis of this trade study. These unused criteria were left in the spreadsheet so that others may insert the criteria weights for their specific application.

Selecton Criteria	Wt	STK 9940B	HP LTO2	IBM LTO2	SDLT 600	Sony SAIT	IBM 3592	STK 9940B	HP LTO2	IBM LTO2	SDLT 600	Sony SAIT	IBM 3592
Design criteria	33	10.0	7.1	7.1	6.5	6.5	7.9	330.0	234.3	234.3	214.5	214.5	260.7
Capacity	10	4.1	4.1	4.1	6.8	10.0	6.0	41.0	41.0	41.0	68.0	100.0	60.0
Media cost/TB		10.0	9.8	9.8	9.7	9.3	7.9	0.0	0.0	0.0	0.0	0.0	0.0
Compatibility	15	1.7	8.3	8.3	10.0	10.0	6.7	25.5	124.5	124.5	150.0	150.0	100.5
Transfer rate		8.2	8.4	8.9	10.0	9.0	6.9	0.0	0.0	0.0	0.0	0.0	0.0
Drive cost		2.4	10.0	6.6	7.7	3.8	1.2	0.0	0.0	0.0	0.0	0.0	0.0
Vendor analyses		4.1	5.8	7.3	9.6	7.6	10.0	0.0	0.0	0.0	0.0	0.0	0.0
Scenario cost	17	4.7	10.0	9.2	9.4	7.6	4.1	79.9	170.0	156.4	159.8	129.2	69.7
Total Weighted Score								**476.4**	**569.8**	**556.2**	**592.3**	**593.7**	**490.9**

Table 4-1 Weighted Decision Matrix

4.2 Conclusions and notes

Although SAIT scored well in this analysis, there are some concerns that can only be addressed through extensive testing. At this time there is no compelling reason to adopt a new standard archive device, though further investigation of SAIT technology is advised prior to future consideration.

- The Sony SAIT-1, the Quantum SDLT 600 and the IBM 3592 were not tested for this report. Performance and capacity figures were based on vendor benchmarks where available, or based on drive specifications combined with past performance (percentage of the claimed specs that were achievable in the past).

- The most heavily weighted criterion is Design Criteria (reliability) and 9940B leads in this crucial category. These findings do not imply that the SDLT 600, SAIT, 3592, and LTO2 are unreliable. When multiple copies of a dataset are maintained, it becomes acceptable to trade cost and performance for reliability.

- The Design Criteria score of the SAIT was difficult to rate due to the use of helical scan technology rather than serpentine. There are no serpentine passes, but past experience and a working knowledge of helical scan technology led to it being rated equivalent to a serpentine drive with a high number of passes.

- As any drive saturates the market, media costs drop.

4.3 Recommendations

1. It is advised that the USGS obtain access to a Quantum SDLT 600 and a Sony SAIT-1 for testing. This access may be remote.

2. It is advised that EDC obtain access to an LTO3 drive when they become available, which is anticipated before this study is revisited.

3. Presuming that the results of further testing prove SAIT and LTO3 to be capable, it is advised that a SAIT drive and an LTO3 drive be procured for more extensive testing.

4. Several of the drives and media have onboard memory which may be used to monitor performance and errors. A future revision of this study should rate these capabilities.

5. In order to reduce risk, the USGS should continue the strategy of storing datasets on multiple technologies. An example of this would be to store a working copy of a dataset on nearline 9940B, and offline/offsite copies on LTO2. This strategy partially mitigates the risks of one or the other technology failing or being retired prematurely.

6. The USGS should adopt a policy of periodically testing archive tapes for readability. This testing should not be extensive enough to incur undue wear on the media, but should be frequent enough to provide an opportunity to detect deteriorating media.

7. It is advised that media be migrated to a new technology at least every 5 years. While most tape technologies can reliably store data for much longer periods, after 5 years the transfer rates and densities that once were leading edge will become problematic, and drives will become difficult to maintain.

8. The USGS should plan to update this trade study periodically.

9. When writing archive tapes, the tapes should be verified on a second drive. This will help identify any drive incompatibility.

10. Each time this study is revisited, it is likely that the leading technology will change. This does not indicate that the USGS should change offline tape technologies frequently. There is benefit in staying with a given technology for several years, even if it is not the leading technology continuously. Specifically, following the 2002 revision, the USGS did not adopt 9940B as the offline technology of choice, due mostly to cost. Following the 2003 revision, the USGS did move quickly to adopt LTO2 as the offline technology of choice – and this has worked well. Following this 2004 revision, it is not advised that the USGS drop LTO2 in favor of SAIT. There currently is no compelling reason to abandon LTO technology, though SAIT warrants further analysis.

Abbreviations and Acronyms

AIT	Advanced Intelligent Tape
CD-ROM	Compact Disc - Read Only Memory
CERN	Conseil European pour la Recherché Nucleaire (European Laboratory for Particle Physics; Geneva, Switzerland)
CRC	Cyclic Redundancy Check
DCT	Digital Cassette tape
DLT	Digital Linear Tape
DVD	Digital Video Disc
EDC	EROS Data Center
EROS	Earth Resources Observation Systems
FYyy	Fiscal Year yy
GB	Gigabytes
HDT	High Density Tape
HP	Hewlett Packard
IBM	International Business Machines
IRIG	InteRange Instrumentation Group (timecode format)
LACS	Landsat Archive Conversion System
LP-DAAC	Land Processes Distributed Active Archive Center
LTO	Linear Tape Open
MB	Megabyte
MODIS	Moderate Resolution Imaging Spectroradiometer
MSS	Multi-spectral Scanner
QIC	Quarter-inch Cartridge
SAIC	Science Applications International Corporation
SAIT	Super Advanced Intelligent Tape
SDLT	Super Digital Linear Tape
TB	Terabytes
WORM	Write Once, Read Many
TM	Thematic Mapper
TMACS	TMMSS Archive Conversion System
USGS	United States Geological Survey